Competing risk modelling in reliability

Tim Bedford
Department of Management Science
Strathclyde University
Glasgow
UK
tim@mansci.strath.ac.uk

Abstract

We give a survey of some of the applications of competing risk models to reliability problems, focussing particularly on some of the models that have been proposed to "remove" the effect of preventive maintenance.

1 Introduction

The competing risk problem arises quite naturally in the reliability context. Maintenance logs often track the history of events occurring at a particular socket. The events can be failure mode specific, incipient failures, maintenance actions, etc. Where the cost of critical failure is large, the maintenance policy will ensure that the whole system is as good as new. Hence we can regard the data as arising from a renewal process in which we only see the "first" possible event occurring after renewal and we know what that event is. The different events can be regarded as competing risks. The "competing risk problem" is that we cannot identify the marginal distributions of the time to each event. Such information is, at least implicitly, used in reliability databases such as the CCPS and EIREDA generic databases.

The presentation gives an overview of competing risk work, specifically in the reliability context. In this sense it complements other more general work (for example Crowder's recent book (Crowder 2001)). In this short written version we shall only give a partial summary.

2 Competing risk framework

In general there may be several different processes going on that could remove a component from service. Hence the time to next removal, Y is the minimum of a number of different potential event times $Y = \min(X_1, \ldots, X_n)$. For simplicity lets just consider a single non-failure event, for example unscheduled preventive maintenance. Hence there is a failure time X which is the time (from the previous service removal) that the equipment would fail, and a single PM time Z which is the time at which the equipment would be preventively maintained. We only observe the smallest of the two variables but also observe which one it is, that is, we know whether we have observed a failure or a PM (we assume here that the two events could not occur simultaneously). Hence the observable data is of the form $Y = (\min(X, Z), 1_{X<Z})$. It would clearly be interesting to know about the distribution of X, that is, the behaviour of the system with the maintenance effect removed. However, we cannot observe X directly. The observations we have allow us only to estimate the *sub-distribution function* $F^*_X(t) = P(X \leq t, X < Z)$. The sub-distribution function increases to the value $P(X < Z)$ as t increases. We often talk about the *sub-survivor function* $S^*_X(t) = P(X > t, X < Z)$, which is equal to $P(X < Z) - F^*_X(t)$. The normalised subsurvivor function is the quantity $S^*_X(t)/S^*_X(0)$ normalised to be equal to 1 at $t = 0$. The final important quantity that can be estimated directly from observable data is the probability of a censor after time t, $\Phi(t) = P(Z < X|Z, X > t)$. The shapes of these functions can play a role in model selection, as we shall see later.

Figure 1 shows why many different models are compatible with the competing risk data. It shows the $X - Z$ plane and the diagonal $x = z$. We can estimate the probability of events such as $X < Z, t_1 < X \leq t_2$, which corresponds to the left, vertical, hatched region in the figure, and the probability of events such as

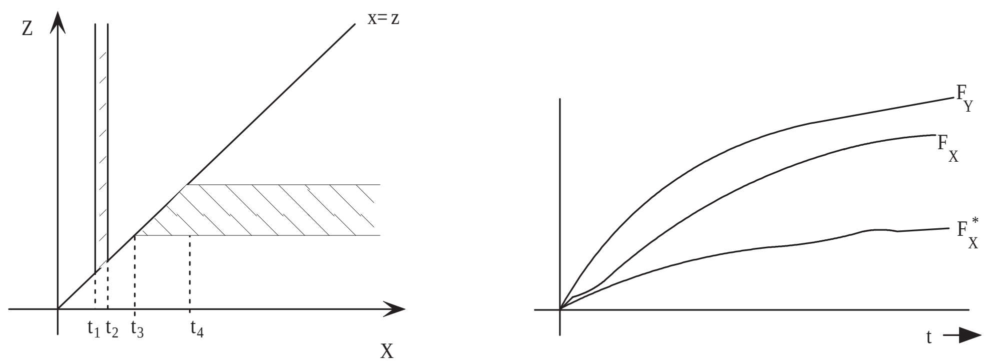

Figure 1: (L) Events whose probabilities can be estimated by competing risk data; (R)Bounding the marginal distribution function

$Z < X, t_3 < Z \leq t_4$, which corresponds to the right, horizontal, hatched region. We are able to estimate the probability of any such region, but we cannot estimate how the probability is distributed *along* such a region. Now we can see why the distribution of X is not identifiable: Varying the mass along the horizontal region changes the distribution of X without changing the distribution of the observable quantities.

3 Model-free bounds

One approach is to try to give bounds on the marginal distribution without making assumptions on the distribution of (X, Z). The Peterson bounds (Peterson 1976) are pointwise upper and lower bounds on the value of the marginal distribution function. They say that, firstly, for any $t \geq 0$ we have $F_X^*(t) \leq F_X(t) \leq F_{\min(X,Z)}(t)$. Furthermore, for any $t \geq 0$ and any u satisfying $F_X^*(t) < u < F_{\min(X,Z)}(t)$, there is a joint distribution on (X, Z) with the given subdistribution functions such that $u = F_X(t)$. See Figure 1.

Functional bounds were given in (Crowder 1991) and improved in (Bedford and Meilijson 1997) to give a complete characterisation of the marginals consistent with given subdistribution functions. Bedford and Meilijson also used the characterisation to produce a statistical test based on the Kolmogorov-Smirnov statistic in which a hypothesised marginal distribution can be tested against available data. The functional bounds say essentially that the gap between each marginal distribution function and its lower bound, should be a non decreasing function t.

The functional bounds can be used to show that the assumption of independence tends to give an optimistic assessment of the marginal of X. Commercial reliability databases appear to assume that the component lifetime is exponential and that the censoring is independent. What happens if the censoring is not independent? It was shown in (Bedford and Meilijson 1995) that any other dependence structure would have given a *higher* estimate of the failure rate: the "independent" failure rate is the lower endpoint of the interval of constant failure rates compatible with the competing risk data.

4 Competing risk models

The models discussed here have been chosen with a view to applications in reliabilty. However, the literature on competing risks is wide, and we refer to (Crowder 2001) for an overview of the general literature.

4.1 Independent competing risks

The model in which X and Z are independent is the most straightforward one, although often implausible for the reasons given above. Estimation is by using the Kaplan-Meier estimator (Crowder 2001).

4.2 Dependent copula model

The sensitivity of predicted lifetime to the assumptions made about dependency between PM and failure time was investigated in (Bunea and Bedford 2002) using a family of dependent copulae.

To illustrate this in an optimisation context, the interpretation given was that of choosing an age-replacement maintenance policy. Existing data, corresponding to failure and/or unscheduled PM events, is taken as input. In order to apply the age replacement maintenance model we need the lifetime distribution of the equipment. Hence it is necessary to "remove" the effect of the unscheduled PM from the lifetime data.

The objective was to show what the costs of assuming the wrong model would be, if one was trying to optimize an age replacement policy. The conclusion of this paper is that the costs of applying the wrong model can indeed be very substantial. The costs arise because the age replacement interval is incorrectly set for the actual lifetime distribution when the lifetime distribution has been incorrectly estimated using false assumptions about the form of censoring.

4.3 Random clipping

This is not really a competing risk model, but is sufficiently close to be included here. The idea, due to Cooke, is that the component life time is exponential, and that the equipment emits a warning at some time before the end of life. That warning period is independent of the lifetime of the equipment, although we only see those warnings that occur while the equipment is on use. The observable data here is the time at which the warning occurs. An application of the memoryless property of the exponential distribution shows that the observable data (the warning times) has the same distribution as the underlying failure time. Hence we can estimate the MTBF just by the mean of the the warning times data. This model, and the following one, is discussed further in (Bedford and Cooke 2001).

4.4 Random signs

This model (Cooke 1993) uses the idea that the time at which PM might occur is related to the time of failure. The PM (censoring time) Z is equal to the failure time X plus an a random quantity ξ, $Z = X - \xi$. Now, while ξ might not be statistically independent of X, its sign is. In other words PM is trying to be effective and to occur round about the time of the failure, but might miss the failure and occur too late. The chance of failure or PM is independent of the time at which the failure would occur. This is quite a plausible model, but is not always compatible with the data. Indeed, Cooke has shown that the model is consistent with the distribution of observable data if and only if the normalised subsurvivor functions are ordered. In other words, this model can be applied if and only if the normalised sub-survivor function for X always lies above that for Z.

4.5 LBL model

This model, proposed in (Langseth and Lindqvist 2003), develops a variant of the random signs model in which the likelihood of an early (that is, before failure) intervention by the maintainer is proportional to the unconditional failure intensity for the component. Maintenance is possibly imperfect in this model. The model is identifiable.

4.6 Mixed exponential model

A new model capturing a class of competing risk data not previously covered by the above was presented in (Bunea, Cooke, and Lindqvist 2002). The underlying model is that X is drawn from a mixture of two exponential distributions, while Z is also exponential and independent of X. This is therefore a special case of the independent competing risks model, but in a very specific parametric setting. Important features of this model that differ from the previous models are: (1) The normalized subdistribution functions are mixtures of exponential distribution functions, (2) The function $\Phi(t)$ increases continuously as a function of t. This model was developed for an application to OREDA data in which these phenomena were observed.

4.7 Delay Time model

The Delay Time model (Christer 2002) is well known within the maintenance community. Here the two times X and Z are expressed in terms of a warning variable and supplementary times, $X = W + X', Z = W + Z'$, where W, X', Z' are mutually independent life variables. As above we observe the minimum of X and Z.

In the case that these variables are all exponential it is fairly straightforward (see for example Section 9.6.4 of (Bedford and Cooke 2001))to see that (a) The normalized subdistribution functions are equal and are exponential distribution functions, (b) The function $\Phi(t)$ is constant as a function of t.

Here we do NOT want to assume that all the variables are exponential and we give a different interpretation to the model. A piece of equipment gives a warning at a time W and is subject to random inspection, with inspections occurring following a Poisson process with parameter λ. Preventive maintenance will be carried out if the inspection occurs after the warning has been given and before the equipment fails. The time to failure after the warning has been given is given by a variable X', which is independent of W. The time to inspection after the warning is exponentially distributed with parameter λ. Hence the time to PM after the warning has been given is an independent exponentially distributed variable Z'. This shows that the model we have discussed above is appropriate for the practical inspection/failure problem, where the only restriction is that Z' is an exponential variable. Using a result from (Bagai and Rao 1992) we can show that this model is identifiable, although we do not have good estimators for it.

5 Conclusion

Dependent competing risk models are increasing being developed to allow the analysis of reliability data. Because of the competing risk problem, non-testable assumptions are required, but it is hoped that appropriate and plausible models can be built to support decision-making

Acknowledgements

I would like to thank my competing risk collaborators over the last few years, in particular Cornel Bunea, Bo Lindqvist, Isaac Meilijson, and especially Roger Cooke who introduced me to the area.

References

Bagai, I. and B. P. Rao (1992). Analysis of survival data with two dependent competing risks. *Biometrika 34*, 801–814.

Bedford, T. and R. Cooke (2001). *Probabilistic Risk Analysis: Foundations and Methods*. CUP.

Bedford, T. and I. Meilijson (1995). The marginal distributions of lifetime variables which right censor each other. In H. Koul and J. Deshpande (Eds.), *IMS Lecture Notes Monograph Series, 27*.

Bedford, T. and I. Meilijson (1997). A characterisation of marginal distributions of (possibly dependent) lifetime variables which right censor each other. *Annals of Statistics 25*, 1622–1645.

Bunea, C. and T. Bedford (2002). The effect of model uncertainty on maintenance optimization. *IEEE Transactions in Reliability 51*, 486 –493.

Bunea, C., R. Cooke, and B. Lindqvist (2002). Competing risk perspective over reliability databases. In H. Langseth and B. Lindqvist (Eds.), *Proceedings of Mathematical Methods in Reliability*. NTNU Press.

Christer, A. (2002). *Stochastic Models in Reliability and Maintenance*, Chapter A review of delay time analysis for modelling plant maintenance. Springer.

Cooke, R. (1993). The total time on test statistic and age-dependent censoring. *Stat. and Prob. Let. 18*.

Crowder, M. (1991). On the identifiability crisis in competing risks analysis. *Scand. J. Statist. 18*, 223–233.

Crowder, M. (2001). *Classical competing risks*. Chapman and Hall/CRC.

Langseth, H. and B. Lindqvist (2003). A maintenance model for components exposed to several failure mechanisms and imperfect repair. In B. L. . K. Doksum (Ed.), *Mathematical and Statistical Methods in Reliability*, pp. 415–430. World Scientific Publishing.

Peterson, A. (1976). Bounds for a joint distribution function with fixed subdistribution functions: Application to competing risks. *Proc. Nat. Acad. Sci. USA 73*, 11–13.